Fundamental Limits of Cache-Aided Wireless BC: Interplay of Coded-Caching and CSIT Feedback

Jingjing Zhang and Petros Elia

arXiv:1511.03961v2 [cs.IT] 25 Apr 2016

***Abstract*—Building on the recent coded-caching breakthrough by Maddah-Ali and Niesen, the work here considers the K-user cache-aided wireless multi-antenna (MISO) symmetric broadcast channel (BC) with random fading and imperfect feedback, and analyzes the throughput performance as a function of feedback statistics and cache size. In this setting, our work identifies the optimal cache-aided degrees-of-freedom (DoF) within a factor of 4, by identifying near-optimal schemes that exploit the new synergy between coded caching and delayed CSIT, as well as by exploiting the unexplored interplay between caching and feedback-quality.**

The derived limits interestingly reveal that — the combination of imperfect quality current CSIT, delayed CSIT, and coded caching, guarantees that — the DoF gains have an initial offset defined by the quality of current CSIT, and then that the additional gains attributed to coded caching are exponential, in the sense that any linear decrease in the required DoF performance, allows for an exponential reduction in the required cache size.

I. Introduction

Recent work by [1] explored — for the single-stream broadcast setting — how careful caching of content at the receivers, and proper encoding across different users' requested data, can allow for higher communication rates. The key idea was to use coding in order to create multicast opportunities, even if the different users requested different data content. This *coded caching* approach — which went beyond storing popular content closer to the user — involved two phases; the placement phase (during off peak hours) and the delivery phase (during peak hours). During the placement phase, content that was predicted to be popular (a library of commonly requested files), was coded and placed across user's caches. During the delivery phase — which started when users requested specific files from the predicted library of files — the transmitter encoded across different users' requested data content, taking into consideration the requests and the existing cache contents. This approach — which translated to efficient interference removal gains that were termed as 'coded-caching gains' — was shown in [1] to provide substantial performance improvement that far exceeded the 'local' caching gains from the aforementioned traditional 'data push' methods that only pre-store content at local caches.

Our interest here is to explore coded caching, not in the original single-stream setting in [1], but rather in the feedback-aided multi-antenna wireless BC. This wireless and multi-antenna element now automatically brings to the fore a largely unexplored and involved relationship between coded caching and CSIT-type feedback quality. This relationship carries particular importance because both CSIT and coded caching are powerful and crucial ingredients in handling interference, because they are both hard to implement individually, and because their utility is affected by one another (often adversely, as we will see). Our work tries to understand how CSIT and caching resources jointly improve performance, as well as tries to shed some light on the interplay between coded caching and feedback.

1) Motivation for the current work: A main motivation in [1] and in subsequent works, was to employ coded caching to remove interference. Naturally, in wireless networks, the ability to remove interference is very much linked to the quality and timeliness of the available feedback, and thus any attempt to further our understanding of the role of coded caching in these networks, stands to benefit from understanding the interplay between coded caching and (variable quality) feedback. This joint exposition becomes even more meaningful when we consider the connections that exist between feedback-usefulness and cached side-information at receivers, where principally the more side information receivers have, the less feedback information the transmitter might need.

This approach is also motivated by the fact that feedback is hard to get in a timely manner, and hence is typically far from ideal and perfect. Thus, given the underlying links between the two, perhaps the strongest reason to jointly consider coded caching and feedback, comes from the prospect of using coded caching to alleviate the constant need to gather and distribute CSIT, which — given typical coherence durations — is an intensive task that may have to be repeated hundreds of times per second during the transmission of content. This suggests that content prediction of a predetermined library of files during the night (off peak hours), and a subsequent caching of parts of this library content again during the night, may go beyond boosting performance, and may in fact offer the additional benefit of alleviating the need for prediction, estimation, and communication of CSIT during the day, whenever requested files are from the library. Our idea of exploring the interplay between feedback (timeliness and quality) and coded caching, hence draws directly from this attractive promise that content prediction, once a day, can offer

The authors are with the Mobile Communications Department at EURECOM, Sophia Antipolis, 06410, France (email: jingjing.zhang@eurecom.fr, elia@eurecom.fr). The work of Petros Elia was supported by the European Community's Seventh Framework Programme (FP7/2007-2013) / grant agreement no.318306 (NEWCOM#), and from the ANR Jeunes Chercheurs project ECOLOGICAL-BITS-AND-FLOPS.

An initial version of this paper has been reported as Research Report No. RR-15-307 at EURECOM, August 25, 2015, http://www.eurecom.fr/publication/4723 as well as was uploaded on arxiv in November 2015.

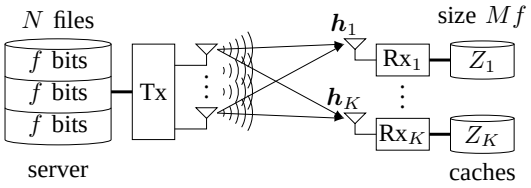

Fig. 1. Cache-aided K-user MISO BC.

repeated and prolonged savings in CSIT.

A. Cache-aided broadcast channel model

1) K-user BC with pre-filled caching: In the symmetric K-user multiple-input single-output (MISO) broadcast channel of interest here, the K-antenna transmitter, communicates to K single-antenna receiving users. The transmitter has access to a library of $N \geq K$ distinct files $W_1, W_2, \ldots, W_N$, each of size $|W_n| = f$ bits. Each user $k \in \{1, 2, \ldots, K\}$ has a cache Z_k, of size $|Z_k| = Mf$ bits, where naturally $M \leq N$. Communication consists of the aforementioned *content placement phase* and the *delivery phase*. During the placement phase — which usually corresponds to communication during off-peak hours — the caches $Z_1, Z_2, \ldots, Z_K$ are pre-filled with content from the N files $\{W_n\}_{n=1}^N$. The delivery phase commences when each user k requests from the transmitter, any *one* file $W_{R_k} \in \{W_n\}_{n=1}^N$, out of the N library files. Each file can be requested with equal probability. Upon notification of the users' requests, the transmitter aims to deliver the (remaining of the) requested files, each to their intended receiver, and the challenge is to do so over a limited (delivery phase) duration T.

For each transmission, the received signals at each user k, will be modeled as

$$y_k = \boldsymbol{h}_k^T \boldsymbol{x} + z_k, \quad k = 1, \ldots, K \tag{1}$$

where $\boldsymbol{x} \in \mathbb{C}^{K\times 1}$ denotes the transmitted vector satisfying a power constraint $\mathbb{E}(||\boldsymbol{x}||^2) \leq P$, where $\boldsymbol{h}_k \in \mathbb{C}^{K\times 1}$ denotes the channel of user k in the form of the random vector of fading coefficients that can change in time and space, and where z_k represents unit-power AWGN noise at receiver k. At the end of the delivery phase, each receiving user k combines the received signal observations y_k — accumulated during the delivery phase — with the fixed information in their respective cache Z_k, to reconstruct their desired file W_{R_k}.

B. Coded caching and CSIT-type feedback

Communication also takes place in the presence of channel state information at the transmitter. CSIT-type feedback is crucial in handling interference, and can thus substantially reduce the resulting duration T of the delivery phase. This CSIT is typically of imperfect-quality as it is hard to obtain in a timely and reliable manner. In the high-SNR (high P) regime of interest, this current-CSIT quality is concisely represented in the form of the normalized quality exponent [2] [3]

$$\alpha := -\lim_{P\to\infty} \frac{\log \mathbb{E}[||\boldsymbol{h}_k - \hat{\boldsymbol{h}}_k||^2]}{\log P}, \ k \in \{1, \ldots, K\} \tag{2}$$

where $\boldsymbol{h}_k - \hat{\boldsymbol{h}}_k$ denotes the estimation error between the current CSIT estimate $\hat{\boldsymbol{h}}_k$ and the estimated channel $\boldsymbol{h}_k$. The range of interest[1] is $\alpha \in [0, 1]$. We also assume availability of delayed CSIT (as in for example [6], as well as in a variety of subsequent works [2], [3], [7]–[14], see also [15]–[18] as well as [19]–[21]) where now the delayed estimates of any channel, can be received without error but with arbitrary delay, even if this delay renders this CSIT completely obsolete. As it is argued in [2], this mixed CSI model (partial current CSIT, and delayed CSIT) nicely captures different realistic settings that might involve channel correlations and an ability to improve CSI as time progresses. This same CSI model is particularly well suited for our caching-related setting here, because it explicitly reflects two key ingredients that are directly intertwined with coded caching; namely, feedback timeliness and feedback quality.

In terms of caching, we will consider the normalized

$$\gamma := \frac{M}{N} \tag{3}$$

as well as the cumulative

$$\Gamma := \frac{KM}{N} = K\gamma. \tag{4}$$

The latter simply means that the sum of the sizes of the caches across all users, is Γ times the volume of the N-file library. As in [1], we will consider the case where $\Gamma = \{1, 2, \cdots K\}$.

a) Intuitive links between α and γ: As we will see, α is not only linked to the performance — where a higher α allows for better interference management and higher performance over the wireless delivery link — but is also linked to caching; after all, the bigger the γ, the more side information the receivers have, the less interference one needs to handle (at least in symmetric systems), and the smaller the α that is potentially needed to steer interference. This means that principally, a higher γ implies that more common information needs to be transmitted, which may (in some cases) diminish the utility of feedback which primarily aims to facilitate the opposite which is the transmission of private information. It is for example easy to see (we will see this later) that in the presence of $\Gamma = K - 1$, there is no need for CSIT in order to achieve the optimal performance.

C. Measures of performance in current work

As in [1], the measure of performance here is the duration T — in time slots, per file served per user — needed to complete the delivery process, *for any request*. The wireless link capabilities, and the time scale, are normalized such that one time slot corresponds to the optimal amount of time it would take to communicate a single file to a single receiver, had there been no caching and no interference. As a result, in the high P setting of interest — where the capacity of a single-user MISO channel scales as $\log_2(P)$ — we proceed to set

$$f = \log_2(P) \tag{5}$$

[1]In the high SNR regime of interest here, $\alpha = 0$ corresponds to having essentially no current CSIT (cf. [4]), while having $\alpha = 1$ corresponds (again in the high SNR regime) to perfect and immediately available CSIT (cf. [5]).

which guarantees that the two measures of performance, here and in [1], are the same and can thus be directly compared[2].

A simple inversion leads to the equivalent measure of the per-user DoF

$$d(\gamma,\alpha)=\frac{1-\gamma}{T} \tag{6}$$

which captures the joint effect of coded caching and feedback[3].

D. Notation and assumptions

1) Notation: We will use the notation $H_n := \sum_{i=1}^{n}\frac{1}{i}$, to represent the n-th harmonic number, and we will use $\epsilon_n := H_n - \log(n)$ to represent its logarithmic approximation error, for some integer n. We remind the reader that ϵ_n decreases with n, and that $\epsilon_\infty := \lim_{n\to\infty} H_n - \log(n)$ is approximately 0.5772. $\mathbb{Z}$ will represent the integers, $\mathbb{Z}^+$ the positive integers, $\mathbb{R}$ the real numbers, $\binom{n}{k}$ the n-choose-k operator, and $\oplus$ the bitwise XOR operation. We will use $[K] := \{1,2,\cdots,K\}$. If ψ is a set, then $|\psi|$ will denote its cardinality. For sets A and B, then $A\backslash B$ denotes the difference set. Complex vectors will be denoted by lower-case bold font. We will use $||\boldsymbol{x}||^2$ to denote the magnitude of a vector $\boldsymbol{x}$ of complex numbers. For a transmitted vector $\boldsymbol{x}$, we will use $\text{dur}(\boldsymbol{x})$ to denote the transmission duration of that vector. For example, having $\text{dur}(\boldsymbol{x})=\frac{1}{10}T$ would simply mean that the transmission of vector $\boldsymbol{x}$ lasts one tenth of the delivery phase. In our high-P setting of interest, we will also use $\doteq$ to denote *exponential equality*, i.e., we will write $g(P)\doteq P^B$ to denote $\lim_{P\to\infty}\frac{\log g(P)}{\log P}=B$. Similarly $\dot{\geq}$ and $\dot{\leq}$ will denote exponential inequalities. Logarithms are of base e, unless we use $\log_2(\cdot)$ which will represent a logarithm of base 2.

2) Main assumptions: Throughout this work, we assume availability of current CSIT with some quality α, of delayed CSIT (D-CSIT), as well as ask that each receiver knows their own channel perfectly. We also adhere to the common convention (see for example [6]) of assuming perfect and global knowledge of delayed channel state information at the receivers (delayed global CSIR), where each receiver must know (with delay) the CSIR of (some of the) other receivers. We will assume that the entries of *each specific* estimation error vector are i.i.d. Gaussian. For the outer (lower) bound to hold, we will make the common assumption that the current channel state must be independent of the previous channel-estimates and estimation errors, *conditioned on the current estimate* (there is no need for the channel to be i.i.d. in time). We will make the assumption that the channel is drawn from a continuous ergodic distribution such that all the channel matrices and all their sub-matrices are full rank almost surely. We also make the soft assumption that the transmitter *during the delivery phase* is aware of the feedback statistics. We note though that, while our main scheme assumes knowledge of α during the caching phase, most results will be the outcome of a simpler scheme that does not require knowledge of α during this caching phase. Removing this assumption entails, for $\alpha>0$, a performance penalty which is small.

E. Prior work

The benefits of coded caching on reducing interference and improving performance, were revealed in the seminal work by Maddah-Ali and Niesen in [1] who considered a caching system where a server is connected to multiple users through a shared link, and designed a novel caching and delivery method that jointly offers a multicast gain that helps mitigate the link load, and which was proven to have a gap from optimal that is at most 12. This work was subsequently generalized in different settings, which included the setting of different cache sizes for which Wang et al. in [26] developed a variant of the algorithm in [1] which achieves a gap of at most 12 from the information theoretic optimal. Other extensions included the work in [27] by Maddah-Ali and Niesen who considered the setting of decentralized caching where the achieved performance was shown to be comparable to that of the centralized case [1], despite the lack of coordination in content placement. For the same original single-stream setting of [1], the work of Ji et al. in [28] considered a scenario where users make multiple requests each, and proposed a scheme that has a gap to optimal that is less than 18. Again for the setting in [1], the work of Ghasemi and Ramamoorthy in [29], derived tighter outer (lower) bounds that improve upon existing bounds, and did so by recasting the bound problem as one of optimally labeling the leaves of a directed tree. Further work can be found in [30] where Wang et al. explored the interesting link between caching and distributed source coding with side information. Interesting conclusions are also drawn in the work of Ajaykrishnan et al. in [31], which revealed that the effectiveness of caching in the single stream case, is diminished when N approaches and exceeds K^2.

Deviating from single-stream error free links, different works have considered the use of coded caching in different wireless networks, without though particular consideration for CSIT feedback quality. For example, work by Huang et al. in [32], considered a cache-aided wireless fading BC where each user experiences a different link quality, and proposed a suboptimal communication scheme that is based on time- and frequency-division and power- and bandwidth-allocation, and which was evaluated using numerical simulations to eventually show that the produced throughput decreases as the number of users increases. Further work by Timo and Wigger in [33] considered an erasure broadcast channel and explored how the cache-aided system efficiency can improve by employing unequal cache sizes that are functions of the different channel qualities. Another work can be found in [34] where Maddah-Ali and Niesen studied the wireless interference channel where each transmitter has a local cache, and showed distinct benefits of coded caching that stem from the fact that content-overlap at the transmitters allows effective interference cancellation.

Different work has also considered the effects of caching in different non-classical channel paradigms. One of the earlier

[2] We note that setting $f=\log_2(P)$ is simply a normalization of choice, and does not carry a 'forced' relationship between SNR and file sizes. The essence of the derived results would remain the same for any other non-trivial normalization.

[3] The DoF measure is designed to exclude the benefits of having some content already available at the receivers (local caching gain), and thus to limit the DoF between 0, and the interference free optimal DoF of 1.

such works that focused on practical wireless network settings, includes the work by Golrezaei et al. in [35], which considered a downlink cellular setting where the base station is assisted by helper nodes that jointly form a wireless distributed caching network (no coded caching) where popular files are cached, resulting in a substantial increase to the allowable number of users by as much as $400-500\%$. In a somewhat related setting, the work in [36] by Perabathini et al. accentuated the energy efficiency gains from caching. Further work by Ji et al. in [37] derived the limits of so-called combination caching networks in which a source is connected to multiple user nodes through a layer of relay nodes, such that each user node with caching is connected to a distinct subset of the relay nodes. Additional work can also be found in [38] where Niesen et al. considered a cache-aided network where each node is randomly located inside a square, and it requests a message that is available in different caches distributed around the square. Further related work on caching can be found in [28], [39]–[43].

Work that combines caching and feedback considerations in wireless networks, has only just recently started. A reference that combines these, can be found in [44] where Deghel et al. considered a MIMO interference channel (IC) with caches at the transmitters. In this setting, whenever the requested data resides within the pre-filled caches, the data-transfer load of the backhaul link is alleviated, thus allowing for these links to be instead used for exchanging CSIT that supports interference alignment. An even more recent concurrent work can be found in [45] where Ghorbel et al. studied the capacity of the cache-enabled broadcast packet erasure channel with ACK/NACK feedback. In this setting, Ghorbel et al. cleverly showed — interestingly also using a retrospective type algorithm, this time by Gatzianas et al. in [46] — how feedback can improve performance by informing the transmitter when to resend the packets that are not received by the intended user and which are received by unintended users, thus allowing for multicast opportunities. The first work that considers the actual interplay between coded caching and CSIT quality, can be found in [47] which considered the easier problem of how the optimal cache-aided performance (with coded caching), can be achieved with reduced quality CSIT.

F. Outline and contributions

In Section II, Lemma 1, we offer a lower bound for the optimal $T^*(\gamma,\alpha)$. Then in Theorem 2 we calculate the achievable $T(\gamma,\alpha)$, for $\Gamma \in \{1,2,\cdots,K\}$, $\alpha \in [0,1]$, and prove it to be less than four times the optimal, thus identifying the optimal $T^*(\gamma,\alpha)$ within a factor of 4. A simpler expression for T (again within a factor of 4 from optimal), and its corresponding per-user DoF, are derived in Theorem 2, while a simple approximation of these is derived in Corollary 2a, where we see that the per-user DoF takes the form $d(\gamma,\alpha) = \alpha + (1-\alpha)\frac{1-\gamma}{\log\frac{1}{\gamma}}$, revealing that even a very small $\gamma = e^{-G}$ can offer a substantial DoF boost $d(\gamma=e^{-G},\alpha) - d(\gamma=0,\alpha) \approx (1-\alpha)\frac{1}{G}$.

In Section III we discuss practical implications. In Corollary 2b we describe the savings in current CSIT that we can have due to coded caching, while in Corollary 2d we quantify the intuition that, in the presence of coded-caching, there is no reason to improve CSIT beyond a certain threshold quality. Furthermore in Section III-B we show how cache-aided communications can utilize a vanishingly-small portion of D-CSIT compared to traditional D-CSIT schemes, simply because caching helps 'skip' the parts of the schemes that require the highest D-CSIT load.

In Section IV we present the caching-and-delivery schemes, which build on the interesting connections between MAT-type retrospective transmission schemes (cf. [6]) and coded caching. The caching part is modified from [1] to essentially *'fold'* (linearly combine) the different users' data into multi-layered blocks, in a way such that the subsequent transmission algorithm (which employs parts of the QMAT algorithm in [48]) is suited to efficiently unfold these. The caching and transmission algorithms are calibrated so that the caching algorithm — which is modified from that in [1] to adapt the caching redundancy to α — creates the same multi-destination delivery problem that is efficiently solved by the last stages of the QMAT scheme. Section VI in the Appendix presents the outer bound proof, and the proof for the gap to optimal.

II. Throughput of cache-aided BC as a function of CSIT quality and caching resources

The following results hold for the (K,M,N,α) cache-aided K-user wireless MISO BC with random fading, $\alpha \in [0,1]$ and $N \geq K$, where $\gamma = \frac{M}{N}$ and $\Gamma = K\gamma$. We begin with an outer bound (lower bound) on the optimal T^*.

Lemma 1: The optimal T^* for the (K,M,N,α) cache-aided K-user MISO BC, is lower bounded as

$$T^*(\gamma,\alpha) \geq \max_{s\in\{1,\ldots,\lfloor\frac{N}{M}\rfloor\}} \frac{1}{(H_s\alpha+1-\alpha)}(H_s - \frac{Ms}{\lfloor\frac{N}{s}\rfloor}). \quad (7)$$

Proof: The proof is presented in Section VI-A and it uses the bound from Lemma 2 whose proof can be found in Section VI-B. ∎

A. Achievable throughput of the cache-aided BC

The following identifies, up to a factor of 4, the optimal T^*, for all $\Gamma \in \{1,2,\cdots,K\}$ (i.e., $M \in \frac{N}{K}\{1,\cdots,K\}$). The result uses the expression

$$\alpha_{b,\eta} = \frac{\eta-\Gamma}{\Gamma(H_K - H_\eta - 1)+\eta},\ \eta = \lceil\Gamma\rceil,\ldots,K-1. \quad (8)$$

Note that the above does not hold for $\Gamma = K$, as this would imply no need for delivery.

Theorem 1: In the (K,M,N,α) cache-aided MISO BC with N files, $K \leq N$ users, $\Gamma \in \{1,2,\cdots,K\}$, and for $\eta = \arg\max_{\eta'\in[\Gamma,K-1]\cap\mathbb{Z}}\{\eta' : \alpha_{b,\eta'} \leq \alpha\}$, then

$$T = \max\{1-\gamma, \frac{(K-\Gamma)(H_K-H_\eta)}{(K-\eta)+\alpha(\eta+K(H_K-H_\eta-1))}\} \quad (9)$$

is achievable and always has a gap-to-optimal that is less than 4, for all α, K. For $\alpha \geq \frac{K(1-\gamma)-1}{(K-1)(1-\gamma)}$, T is optimal.

Proof: The caching and delivery scheme that achieves the above performance is presented in Section IV, while the corresponding gap to optimal is bounded in Section VI-C. ■

The above is achieved with a general scheme whose caching phase is a function of α. We will henceforth consider a special case ($\eta = \Gamma$) of this scheme, which provides similar performance (it again has a gap to optimal that is bounded by 4), simpler expressions, and has the practical advantage that the caching phase need not depend on the CSIT statistics α of the delivery phase. For this case, we can achieve the following performance.

Theorem 2: In the (K, M, N, α) cache-aided MISO BC with $\Gamma \in \{1, 2, \cdots, K\}$,

$$T = \frac{(1-\gamma)(H_K - H_\Gamma)}{\alpha(H_K - H_\Gamma) + (1-\alpha)(1-\gamma)} \tag{10}$$

is achievable and has a gap from optimal

$$\frac{T}{T^*} < 4 \tag{11}$$

that is less than 4, for all α, K. Thus the corresponding per-user DoF takes the form

$$d(\gamma, \alpha) = \alpha + (1-\alpha)\frac{1-\gamma}{H_K - H_\Gamma}. \tag{12}$$

Proof: The scheme that achieves the above performance will be described later on as a special (simpler) case of the scheme corresponding to Theorem 1. The corresponding gap to optimal is bounded in Section VI-C. ■

The following corollary describes the above achievable T, under the logarithmic approximation $H_n \approx \log(n)$. The presented expression is exact in the large K setting[4] where $\frac{H_K - H_\Gamma}{\log(\frac{1}{\gamma})} = 1$.

Corollary 2a: Under the logarithmic approximation $H_n \approx \log(n)$, the derived T takes the form

$$T(\gamma, \alpha) = \frac{(1-\gamma)\log(\frac{1}{\gamma})}{\alpha\log(\frac{1}{\gamma}) + (1-\alpha)(1-\gamma)} \tag{13}$$

and the derived DoF takes the form

$$d(\gamma, \alpha) = \alpha + (1-\alpha)\frac{1-\gamma}{\log\frac{1}{\gamma}}. \tag{14}$$

For the large K setting, what the above suggests is that current CSIT offers an initial DoF boost of $d^*(\gamma = 0, \alpha) = \alpha$ (cf. [48]), which is then supplemented by a DoF gain

$$d(\gamma, \alpha) - d^*(\gamma = 0, \alpha) \to (1-\alpha)\frac{1-\gamma}{\log(\frac{1}{\gamma})}$$

attributed to the synergy between delayed CSIT and caching [5]. These synergistic gains (see also [49]) are accentuated for smaller values of γ, where we see an exponential effect of coded caching, in the sense that now a microscopic $\gamma = e^{-G}$ can offer a substantial DoF boost

$$d(\gamma = e^{-G}, \alpha) - d(\gamma = 0, \alpha) \approx (1-\alpha)\frac{1}{G}. \tag{15}$$

Example 1: In a MISO BC system with $\alpha = 0$, K antennas and K users, in the absence of caching, the optimal per-user DoF is $d^*(\gamma = 0, \alpha = 0) = 1/H_K$ (cf. [6]) which vanishes to zero as K increases. A DoF of $1/4$ can be guaranteed with $\gamma \approx \frac{1}{50}$ for all K, a DoF of $1/7$ with $\gamma \approx \frac{1}{1000}$, and a DoF of $1/11.7$ can be achieved with $\gamma \approx 10^{-5}$, again for all K.

a) Interplay between CSIT quality and coded caching in the symmetric MISO BC: The derived form in (12) (and its approximation in (14)) nicely capture the synergistic as well as competing nature of feedback and coded caching. It is easy to see for example that the effect from coded-caching, reduces with α and is proportional to $1-\alpha$. This reflects the fact that in the symmetric MISO BC, feedback supports broadcasting by separating data streams, thus diminishing multi-casting by reducing the number of common streams. In the extreme case when $\alpha = 1$, we see — again for the symmetric MISO BC — that the caching gains are limited to local caching gains[6].

III. Cache-aided CSIT reductions

We proceed to explore how coded caching can alleviate the need for CSIT.

A. Cache-aided CSIT gains

To capture the cache-aided reductions on the CSIT load, let us consider

$$\bar{\alpha}(\gamma, \alpha) := \arg\min_{\alpha'} \{\alpha' : (1-\gamma)T^*(\gamma = 0, \alpha') \leq T(\gamma, \alpha)\} \tag{16}$$

which is derived below in the form

$$\bar{\alpha}(\gamma, \alpha) = \alpha + \delta_\alpha(\gamma, \alpha)$$

for some $\delta_\alpha(\gamma, \alpha)$ that can be seen as the *CSIT reduction due to caching* (from $\bar{\alpha}(\gamma, \alpha)$ to the operational α).

Corollary 2b: In the (K, M, N, α) cache-aided MISO BC, then

$$\bar{\alpha}(\gamma, \alpha) = \alpha + \frac{(1-\alpha)(H_{K\gamma} - \gamma H_K)}{(H_K - 1)(H_K - H_{K\gamma})} \tag{17}$$

is achievable, and implies a cache-aided CSIT reduction

$$\delta_\alpha(\gamma, \alpha) = \frac{(1-\alpha)(H_{K\gamma} - \gamma H_K)}{(H_K - 1)(H_K - H_{K\gamma})}.$$

Proof: The proof is direct from Theorem 2. ■

The above is made more insightful in the large K regime, for which we have the following.

[4]For large K, this approximation $\frac{H_K - H_\Gamma}{\log(\frac{1}{\gamma})} = 1$ is tight for any *fixed* γ.

[5]We note that these interference-removal gains, particularly in the large K regime, are not a result of extra performance boost directly from D-CSIT, because in the large K setting, this latter performance boost is negligible (vanishes to zero) without caching.

[6]This conclusion is general (and not dependent on the specific schemes), because the used schemes are optimal for $\alpha = 1$. The statement holds because we can simply uniformly cache a fraction γ of each file in each cache, and upon request, use perfect-CSIT to zero-force the remaining requested information, to achieve the optimal $T^*(\gamma, \alpha = 1) = 1 - \gamma$, which leaves us with local (data push) caching gains only.

Corollary 2c: In the (K, M, N, α) cache-aided MISO BC, then

$$\bar{\alpha}(\gamma, \alpha) = \alpha + (1-\alpha)\frac{1-\gamma}{\log(\frac{1}{\gamma})} \tag{18}$$

which implies CSIT reductions of

$$\delta_\alpha(\gamma, \alpha) = (1-\alpha)d(\gamma, \alpha = 0) = (1-\alpha)\frac{1-\gamma}{\log(\frac{1}{\gamma})}.$$

Proof: The proof is direct from the definition of $\bar{\alpha}(\gamma, \alpha)$ and from Theorem 2. ∎

Furthermore we have the following which quantifies the intuition that, in the presence of coded-caching, there is no reason to improve CSIT beyond a certain threshold quality. The following uses the definition in (8), and it holds for all K.

Corollary 2d: For any $\Gamma \in \{1, \ldots, K\}$, then

$$T^*(\gamma, \alpha) = T^*(\gamma, \alpha = 1) = 1 - \gamma \tag{19}$$

holds for any

$$\alpha \geq \alpha_{b,K-1} = \frac{K(1-\gamma)-1}{(K-1)(1-\gamma)} \tag{20}$$

which reveals that CSIT quality $\alpha = \alpha_{b,K-1}$ is the maximum needed, as it already offers the same optimal performance $T^*(\gamma, \alpha = 1)$ that would be achieved if CSIT was perfect.

Proof: This is seen directly from Theorem 1 after noting that the achievable T matches $T^*(\gamma, \alpha = 1) = 1 - \gamma$. ∎

b) How much caching is needed to partially substitute current CSIT with delayed CSIT (using coded caching to 'buffer' CSI): As we have seen, in addition to offering substantial DoF gains, the synergy between feedback and caching can also be applied to reduce the burden of acquiring current CSIT. What the above results suggest is that a modest γ can allow a BC system with D-CSIT to approach the performance attributed to current CSIT, thus allowing us to partially substitute current with delayed CSIT, which can be interpreted as an ability to buffer CSI. A simple calculation — for the large-K regime — can tell us that

$$\gamma_\alpha^{'} := \arg\min_{\gamma^{'}} \{\gamma^{'} : d(\gamma^{'}, \alpha = 0) \geq d^*(\gamma = 0, \alpha)\} = e^{-1/\alpha}$$

which means that $\gamma_\alpha^{'} = e^{-1/\alpha}$ suffices to achieve — in conjunction with delayed CSIT — the optimal DoF performance $d^*(\gamma = 0, \alpha)$ associated to a system with delayed CSIT and α-quality current CSIT.

Example 2: Let K be very large, and consider a BC system with delayed CSIT and α-quality current CSIT, where $\alpha = 1/5$. Then $\gamma_{\alpha=1/5}^{'} = e^{-5} = 0.0067 \approx 1/150$ which means that

$$d^*(\gamma = 0.0067, \alpha = 0) \geq d^*(\gamma = 0, \alpha = 1/5)$$

which says that the same high-K per-user DoF performance $d^*(\gamma = 0, \alpha = 1/5)$, can be achieved by substituting all current CSIT with coded caching employing $\gamma \approx 1/150$.

B. Vanishing fraction of delayed CSIT

In the following we briefly explore how caching allows for a reduced D-CSIT load. We do so for the case of $\alpha = 0$.

When $\alpha = 0$, the delivery scheme which we describe in Section IV, draws directly from the MAT scheme [6]. This scheme can have up to K phases which are of decreasing time duration and which use a decreasing number of transmit antennas. Essentially each phase is lighter than the previous one, in terms of implementation difficulty. What we will see is that caching will allow us to bypass the first Γ phases, which are the longest and most intensive, leaving us with the remaining $K - \Gamma$ communication phases that are easier to support with delayed feedback because they involve fewer transmissions, with fewer transmit antennas and to fewer users, and thus involve fewer D-CSIT scalars that must be communicated.

In brief — after normalization to account for the condition that each user receives a total of $\log_2(P)$ bits of data — each phase $j = \Gamma+1, \Gamma+2, \ldots, K$ will have a *normalized* duration $T_j = \frac{1}{j}$. During each phase j, we will need to send D-CSIT that describes the channel vectors for $K-j$ users, and during this same phase the transmitted vectors will have support $K-j+1$ because only $K-j+1$ transmit antennas are active. Thus during phase j, there will be a need to send $T_j(K-j+1)(K-j) = \frac{1}{j}(K-j+1)(K-j)$ D-CSIT scalars, and thus a need to send D-CSIT for up to a total of

$$\begin{aligned} L(\Gamma) &= \sum_{j=\Gamma+1}^{K} \frac{1}{j}(K-j+1)(K-j) \\ &= (K^2+K)(H_K - H_\Gamma) - \frac{K(1-\gamma)(3K-K\gamma-1)}{2} \end{aligned}$$

channel scalars, while in the absence of caching (corresponding to $\Gamma = 0$), we will have to send D-CSIT on

$$\begin{aligned} L(\Gamma = 0) &= \sum_{j=1}^{K} \frac{1}{j}(K-j+1)(K-j) \\ &= (K^2+K)H_K - \frac{3K^2}{2} + \frac{K}{2} \end{aligned}$$

channel scalars.

To reflect the frequency of having to gather D-CSIT, and to provide a fair comparison between different schemes of different performance that manage to convey different amounts of actual data to the users, we consider the measure $Q(\Gamma)$ that normalizes the above number $L(\Gamma)$ of full D-CSIT scalars, by the coherence period T_c and by the total number of full data symbols sent. In our case, under the assumption that each user receives a total of $\log_2(P)$ bits, the total number of full data symbols sent is K, and thus we have

$$\begin{aligned} Q(\Gamma) &= \frac{L(\Gamma)}{T_c K} \\ &= \frac{(K^2+K)(H_K - H_\Gamma) - \frac{K(1-\gamma)(3K-K\gamma-1)}{2}}{T_c K} \end{aligned}$$

while without caching, we have

$$Q(\Gamma = 0) = \frac{L(\Gamma)}{T_c K} = \frac{(K+1)H_K - \frac{3}{2}K + \frac{1}{2}}{T_c}.$$

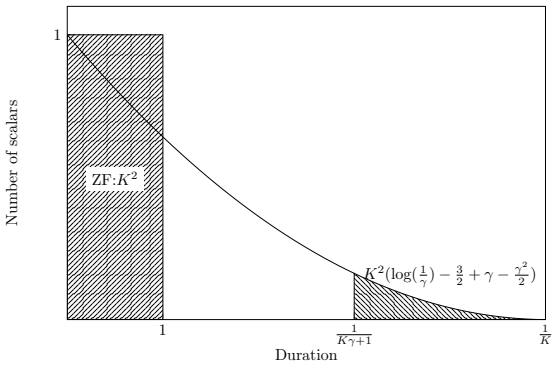

Fig. 2. Illustration of the vanishing fraction of D-CSIT cost, due to caching.

Consequently we see that in the large K limit,

$$Q(\Gamma) \to \frac{K\big(\log(\frac{1}{\gamma}) - \frac{3}{2} + 2\gamma - 2\gamma^2\big)}{T_c}$$

$$Q(\Gamma = 0) \to \frac{1}{T_c} K \log(K)$$

which implies that

$$\lim_{K\to\infty} \frac{Q(\Gamma)}{Q(\Gamma = 0)} = 0$$

which in turn tells us that as K increases, for any fixed γ, caching allows for a substantial reduction (down to a vanishingly small portion) from the original cost of D-CSIT. This is illustrated in Figure 2.

This reduction is important because retrospective delayed-feedback methods suffer from an increased cost of supporting their CSIT requirements (cf. [50]) (albeit at the benefit of allowing substantial delays in the feedback mechanisms); after all, in the presence of perfect CSIT and zero forcing (no caching), the same cost is

$$Q_{ZF} = \frac{K^2}{T_c K} = \frac{K}{T_c}$$

which gives that

$$\lim_{K\to\infty} \frac{Q(\Gamma = 0)}{Q_{ZF}} = \infty$$

which in turn verifies the above claim, and shows that the increase in the cost of supporting the D-CSIT (without caching) can be unbounded compared to ZF methods. On the other hand, we see that

$$\lim_{K\to\infty} \frac{Q(\Gamma)}{Q_{ZF}} = \log(\frac{1}{\gamma} - \frac{3}{2} + 2\gamma - 2\gamma^2)$$

which means that

$$\lim_{K\to\infty} \frac{Q(\Gamma)}{Q_{ZF}} < 1, \ \gamma \geq \frac{1}{10}.$$

One interesting conclusion that comes out of this, is that caching can allow for full substitution of current CSIT (as we have seen above), with a very substantial reduction of the cost of D-CSIT as well, where for $\gamma \geq \frac{1}{10}$ this cost is even less than that of the very efficient ZF, which has to additionally deal though with harder-to-obtain current CSIT. This cost reduction is also translated into a reduction in the cost of disseminating global channel state information at the receivers (global CSIR), where each receiver must now know (again with delay that is allowed to be large) the CSIR of only a fraction of the other receivers.

IV. Cache-aided retrospective communications

We proceed to describe the communication scheme, and in particular the process of placement, folding-and-delivery, and decoding. In the end we calculate the achievable duration T.

The caching part is modified from [1] to *'fold'* (linearly combine) the different users' data into multi-layered blocks, in a way such that the subsequent Q-MAT transmission algorithm (cf. [48]) (specifically the last $K - \eta_\alpha$ ($\eta_\alpha \in \{\Gamma, \dots, K-1\}$) phases of the QMAT algorithm) can efficiently deliver these blocks. Equivalently the algorithms are calibrated so that the caching algorithm creates a multi-destination delivery problem that is the same as that which is efficiently solved by the last stages of the QMAT-type communication scheme. We henceforth remove the subscript in η_α and simply use η, where now the dependence on α is implied.

A. Placement phase

We proceed with the placement phase which modifies on the work of [1] such that when the CSIT quality α increases, the algorithm caches a decreasing portion from each file, but does so with increasing redundancy. The idea is that the higher the α, the more private messages one can deliver directly without the need to multicast, thus allowing for some of the data to remain entirely uncached, which in turn allows for more copies of the same information across different users' caches.

Here each of the N files $W_n, n = 1, 2, \dots, N$ ($|W_n| = f$ bits) in the library, is split into two parts

$$W_n = (W_n^c, W_n^{\overline{c}}) \tag{21}$$

where W_n^c (c for 'cached') will be placed into one or more caches, while the content of $W_n^{\overline{c}}$ ($\overline{c}$ for 'non-cached') will never be cached anywhere, but will instead be communicated — using CSIT — in a manner that causes manageable interference and hence does not necessarily benefit from coded caching. The split is such that

$$|W_n^c| = \frac{KMf}{N\eta} \tag{22}$$

where $\eta \in \{\Gamma, \dots, K-1\}$ is a positive integer, the value of which will be decided later on such that it properly regulates how much to cache from each W_n. Now for any specific η, we equally divide W_n^c into $\binom{K}{\eta}$ subfiles $\{W_{n,\tau}^c\}_{\tau\in\Psi_\eta}$,

$$W_n^c = \{W_{n,\tau}^c\}_{\tau\in\Psi_\eta} \tag{23}$$

where[7]

$$\Psi_\eta := \{\tau \subset [K] \ : \ |\tau| = \eta\} \tag{24}$$

[7]We recall that in the above, τ and $W_{n,\tau}^c$ are sets, thus $|\tau|, |W_{n,\tau}^c|$ denote cardinalities; $|\tau| = \eta$ means that τ has η different elements from $[K]$, while $|W_{n,\tau}^c|$ describes the size of $W_{n,\tau}^c$ in bits.

where each subfile has size

$$|W^c_{n,\tau}| = \frac{KMf}{N\eta\binom{K}{\eta}} = \frac{Mf}{N\binom{K-1}{\eta-1}} \text{ bits.} \quad (25)$$

Now drawing from [1], the caches are filled as follows

$$Z_k = \{W^c_{n,\tau}\}_{n\in[N],\tau\in\Psi^{(k)}_\eta} \quad (26)$$

where

$$\Psi^{(k)}_\eta := \{\tau \in \Psi_\eta \ : \ k \in \tau\}. \quad (27)$$

Hence each subfile $W^c_{n,\tau}$ is stored in Z_k as long as $k \in \tau$, which means that each $W^c_{n,\tau}$ (and thus each part of W^c_n) is repeated η times in the caches. As η increases with α, this means that CSIT allows for a higher redundancy in the caches; instead of content appearing in Γ different caches, it appears in $\eta \geq \Gamma$ caches instead, which will translate into multicast messages that are intended for more receivers.

B. Data folding

At this point, the transmitter becomes aware of the file requests $R_k, k = 1,\dots,K$, and must now deliver each requested file W_{R_k}, by delivering the constituent subfiles $\{W^c_{R_k,\tau}\}_{\tau\in\Psi_\eta\backslash\Psi^{(k)}_\eta}$ as well as $W^{\overline{c}}_{R_k}$, all to the corresponding receiver k. We quickly recall that:

1) subfiles $\{W^c_{R_k,\tau}\}_{\tau\in\Psi^{(k)}_\eta}$ are already in Z_k;
2) subfiles $\{W^c_{R_k,\tau}\}_{\tau\in\Psi_\eta\backslash\Psi^{(k)}_\eta}$ are directly requested by user k, but are not cached in Z_k;
3) subfiles $Z_k\backslash\{W^c_{R_k,\tau}\}_{\tau\in\Psi^{(k)}_\eta} = Z_k\backslash W^c_{R_k}$ are cached in Z_k, are not directly requested by user k, but will be useful in removing interference.

We assume the communication here has duration T. Thus for each k and a chosen η, we split each subfile $W^c_{R_k,\tau}$, $\tau \in \Psi_\eta\backslash\Psi^{(k)}_\eta$ (each of size $|W^c_{R_k,\tau}| = \frac{Mf}{N\binom{K-1}{\eta-1}}$ as we saw in (25)) into

$$W^c_{R_k,\tau} = [W^{c,f}_{R_k,\tau} \quad W^{c,\overline{f}}_{R_k,\tau}] \quad (28)$$

where $W^{c,f}_{R_k,\tau}$ corresponds to information that appears in a cache somewhere and that will be eventually 'folded' (XORed) with other information, whereas $W^{c,\overline{f}}_{R_k,\tau}$ corresponds to information that is cached somewhere but that will not be folded with other information. The split yields

$$|W^{c,\overline{f}}_{R_k,\tau}| = \frac{f\alpha T - f(1-\frac{KM}{N\eta})}{\binom{K-1}{\eta}} \quad (29)$$

where in the above, $f\alpha T$ represents the load for each user without causing interference during the delivery phase, where $f(1-\frac{KM}{N\eta})$ is the amount of uncached information, and where $|W^{c,f}_{R_k,\tau}| = |W^c_{R_k,\tau}| - |W^{c,\overline{f}}_{R_k,\tau}|$.

We proceed to fold cached content, by creating linear combinations (XORs) from $\{W^{c,f}_{R_k,\tau}\}_{\tau\in\Psi_\eta\backslash\Psi^{(k)}_\eta}, \forall k$. We will use $P_{k,k'}(\tau)$ to be the function that replaces inside τ, the entry $k' \in \tau$, with the entry k. As in [1], the idea is that if we deliver

$$W^{c,f}_{R_k,\tau} \oplus (\oplus_{k'\in\tau} \underbrace{W^{c,f}_{R_{k'},P_{k,k'}(\tau)}}_{\in Z_k}) \quad (30)$$

the fact that $W^{c,f}_{R_{k'},P_{k,k'}(\tau)} \in Z_k$, guarantees that receiver k can recover $W^{c,f}_{R_k,\tau}$, while at the same time guarantees that each other user $k' \in \tau$ can recover its own desired subfile $W^{c,f}_{R'_k,P_{k,k'}(\tau)} \notin Z_{k'}, \forall k' \in \tau$.

Hence delivery of each $W^{c,f}_{R_k,\tau} \oplus (\oplus_{k'\in\tau} W^{c,f}_{R_{k'},P_{k,k'}(\tau)})$ of size $|W^{c,f}_{R_k,\tau} \oplus (\oplus_{k'\in\tau} W^{c,f}_{R_{k'},P_{k,k'}(\tau)})| = |W^{c,f}_{R_k,\tau}|$ (cf. (25)), automatically guarantees delivery of $W^{c,f}_{R_{k'},P_{k,k'}(\tau)}$ to each user $k' \in \tau$, i.e., simultaneously delivers a total of $\eta+1$ distinct subfiles (each again of size $|W^{c,f}_{R_{k'},P_{k,k'}(\tau)}| = |W^{c,f}_{R_k,\tau}|$ bits) to $\eta+1$ distinct users. Hence *any*

$$X_\psi := \oplus_{k\in\psi} W^{c,f}_{R_k,\psi\backslash\{k\}}, \psi \in \Psi_{\eta+1} \quad (31)$$

— which is of the same form as in (30), and which is referred to here as an *order-($\eta+1$) folded message* — can similarly deliver to user $k \in \psi$, her requested file $W^{c,f}_{R_k,\psi\backslash k}$, which in turn means that each order-$(\eta+1)$ folded message X_ψ can deliver — with the assistance of the side information in the caches — a distinct, individually requested subfile, to each of the $\eta+1$ users $k \in \psi$ ($\psi \in \Psi_{\eta+1}$).

Thus to satisfy all requests $\{W_{R_k}\backslash Z_k\}_{k=1}^K$, the transmitter must deliver

- uncached messages $W^{\overline{c}}_{R_k}$, $k = 1,\dots,K$
- cached but unfolded messages $\{W^{c,\overline{f}}_{R_k,\psi\backslash\{k\}}\}_{\psi\in\Psi_{\eta+1}}$, $k = 1,\dots,K$
- and the entire set

$$\mathcal{X}_\Psi := \{X_\psi = \oplus_{k\in\psi} W^{c,f}_{R_k,\psi\backslash\{k\}}\}_{\psi\in\Psi_{\eta+1}} \quad (32)$$

consisting of

$$|\mathcal{X}_\Psi| = \binom{K}{\eta+1} \quad (33)$$

folded messages of order-$(\eta+1)$, each of size (cf. (29),(25))

$$|X_\psi| = |W^{c,f}_{R_k,\tau}| = |W^c_{R_k,\tau}| - |W^{c,\overline{f}}_{R_k,\tau}| = \frac{f(1-\gamma-\alpha T)}{\binom{K-1}{\eta}} \text{ (bits).} \quad (34)$$

C. Transmission

The transmission scheme is taken from [48], and each transmission takes the form

$$\boldsymbol{x}_t = \mathbf{G}_{c,t}\boldsymbol{x}_{c,t} + \sum_{k\in\bar{\psi}} \boldsymbol{g}_{k,t}a^*_{k,t} + \sum_{k=1}^K \boldsymbol{g}_{k,t}a_{k,t} \quad (35)$$

where $t \in [0,T]$, where $\boldsymbol{x}_{c,t}$ is a K-length vector for MAT-type symbols, where $a^*_{k,t}$ is the additional auxiliary symbols that carry residual interference (here we 'load' this round with additional requests from the users, for the very first round, $a^*_{k,t} = 0$) (in the above, $\bar{\psi}$ is a set of 'undesired' users). In the above, each unit-norm precoder $\boldsymbol{g}_{k,t}$ for user $k = 1,2,\dots,K$, is simultaneously orthogonal to the CSI of all other channels, i.e.,

$$\hat{\boldsymbol{h}}^T_{k',t}\boldsymbol{g}_{k,t} = 0, \quad \forall k' \in [K]\backslash k. \quad (36)$$

Each precoder $\mathbf{G}_{c,t}$ is defined as $\mathbf{G}_{c,t} = [\boldsymbol{g}_{c,t}, \mathbf{U}_{c,t}]$, where $\boldsymbol{g}_{c,t}$ is simultaneously orthogonal to the channel estimates of the undesired receivers, and $\mathbf{U}_{c,t} \in \mathbb{C}^{K\times(K-1)}$ is a randomly chosen, isotropically distributed unitary matrix[8].

Throughout communication

- we will allocate power such that

$$\mathbb{E}\{|\boldsymbol{x}_{c,t}|_1^2\} \doteq \mathbb{E}\{|a_{k,t}^*|^2\} \doteq P, \qquad (37)$$
$$\mathbb{E}\{|\boldsymbol{x}_{c,t}|_{i\neq 1}^2\} \doteq P^{1-\alpha},\ \mathbb{E}\{|a_{k,t}|^2\} \doteq P^{\alpha}$$

  where $|\boldsymbol{x}_{c,t}|_i, i = 1, 2, \cdots, K$, denotes scalar i in vector $\boldsymbol{x}_{c,t}$, and we will allocate rate such that
- each $\boldsymbol{x}_{c,t}$ carries $f(1-\alpha)$ bits per unit time,
- each $a_{k,t}^*$ carries $f(1-\alpha)$ bits per unit time.
- and each $a_{k,t}$ carries $f\alpha$ bits per unit time.

Remark 1: Recall that instead of employing matrix notation, after normalization, we use the concept of signal duration $\text{dur}(\boldsymbol{x})$ required for the transmission of some vector $\boldsymbol{x}$. We also note that due to time normalization, the time index $t \in [0, T]$, need not be an integer.

For any α, our scheme will be defined by an integer $\eta \in [\Gamma, K-1] \cap \mathbb{Z}$, which will be chosen as

$$\eta = \arg \max_{\eta' \in [\Gamma, K-1]\cap\mathbb{Z}} \{\eta' \ :\ \alpha_{b,\eta'} \leq \alpha\} \qquad (38)$$

for

$$\alpha_{b,\eta} = \frac{\eta - \Gamma}{\Gamma(H_K - H_\eta - 1) + \eta}. \qquad (39)$$

η will define the amount of cached information that will be folded ($\{W_{R_k,\tau}^{c,f}\}_{\tau\in\Psi_\eta\backslash\Psi_\eta^{(k)}}$), and thus also the amount of cached information that will not be folded ($\{W_{R_k,\tau}^{c,\overline{f}}\}_{\tau\in\Psi_\eta\backslash\Psi_\eta^{(k)}}$) and which will be exclusively carried by the different $a_{k,t}$. In all cases,

- all of $\{X_\psi\}_{\psi\in\Psi_{\eta+1}}$ (which are functions of the cached-and-to-be-folded $\{W_{R_k,\tau}^{c,f}\}_{\tau\in\Psi_\eta\backslash\Psi_\eta^{(k)}}$) will be exclusively carried by $\boldsymbol{x}_{c,t},\ t \in [0, T]$, while
- all of the uncached $W_{R_k}^{c}$ (for each $k = 1, \ldots, K$) and all of the cached but unfolded $\{W_{R_k,\tau}^{c,\overline{f}}\}_{\tau\in\Psi_\eta\backslash\Psi_\eta^{(k)}}$ will be exclusively carried by $a_{k,t}, t \in [0, T]$.

Transmission of $\{X_\psi\}_{\psi\in\Psi_{\eta+1}}$: From [48], we know that the transmission relating to $\boldsymbol{x}_{c,t}$ can be treated independently from that of $a_{k,t}$, simply because — as we will further clarify later on — the $a_{k,t}$ do not actually interfere with decoding of $\boldsymbol{x}_{c,t}$, as a result of the scheme, and as a result of the chosen power and rate allocations which jointly adapt to the CSIT quality α. For this reason, we can treat the transmission of $\boldsymbol{x}_{c,t}$ separately.

Hence we first focus on the transmission of $\{X_\psi\}_{\psi\in\Psi_{\eta+1}}$, which will be sent using $\boldsymbol{x}_{c,t},\ t \in [0, T]$ using the last $K-\eta$ phases of the QMAT algorithm in [48] corresponding to having the ZF symbols $a_{k,t}$ set to zero. For ease of notation, we will label these phases starting from phase $\eta+1$ and terminating in phase K. The total duration is the desired T. Each phase $j = \eta+1, \ldots, K$ aims to deliver order-j folded messages (cf. (32)), and will do so gradually: phase j will try to deliver (in addition to other information) $N_j := (K-j+1)\binom{K}{j}$ order-j messages which carry information that has been requested by j users, and in doing so, it will generate $N_{j+1} := j\binom{K}{j+1}$ signals that are linear combinations of received signals from $j+1$ different users, and where these N_{j+1} signals will be conveyed in the next phase $j+1$. During the last phase $j = K$, the transmitter will send fully common symbols that are useful and decoded by all users, thus allowing each user to go back and retroactively decode the information of phase $j = K-1$, which will then be used to decode the information in phase $j = K-2$ and so on, until they reach phase $j = \eta+1$ (first transmission phase) which will complete the task. We proceed to describe these phases. We will use T_j to denote the duration of phase j.

Phase $\eta+1$: In this first phase of duration $T_{\eta+1}$, the information in $\{X_\psi\}_{\psi\in\Psi_{\eta+1}}$ is delivered by $\boldsymbol{x}_{c,t},\ t \in [0, T_{\eta+1}]$, which can also be rewritten in the form of a sequential transmission of shorter-duration K-length vectors

$$\boldsymbol{x}_\psi = [x_{\psi,1}, \ldots, x_{\psi,K-\eta}, 0, \ldots, 0]^T \qquad (40)$$

for different ψ, where each vector $\boldsymbol{x}_\psi$ carries exclusively the information from each X_ψ, and where this information is uniformly split among the $K-\eta$ independent scalar entries $x_{\psi,i},\ i = 1, \ldots, K-\eta$, each carrying

$$\frac{|X_\psi|}{(K-\eta)} = \frac{f(1-\gamma-\alpha T)}{\binom{K-1}{\eta}(K-\eta)} \qquad (41)$$

bits (cf. (34)). Hence, given that the allocated rate for $\boldsymbol{x}_{c,t}$ (and thus the allocated rate for each $\boldsymbol{x}_\psi$) is $(1-\alpha)f$, we have that the duration of each $\boldsymbol{x}_\psi$ is

$$\text{dur}(\boldsymbol{x}_\psi) = \frac{|X_\psi|}{(K-\eta)(1-\alpha)f}. \qquad (42)$$

Given that $|\mathcal{X}_\Psi| = \binom{K}{\eta+1}$, then

$$T_{\eta+1} = \binom{K}{\eta+1}\text{dur}(\boldsymbol{x}_\psi) = \frac{\binom{K}{\eta+1}|X_\psi|}{(K-\eta)(1-\alpha)f}. \qquad (43)$$

After each transmission of $\boldsymbol{x}_\psi$, each user $k \in [K]$ receives (in addition to information originating from $a_{k,t}$ which will be treated as noise and thus neglected for now), a linear combination $L_{\psi,k}$ of the transmitted $K-\eta$ symbols $x_{\psi,1}, x_{\psi,2}, \ldots, x_{\psi,K-\eta}$. Next the transmitter will send an additional $K-\eta-1$ signals $L_{\psi,k'},\ k' \in [K]\backslash\psi$ (linear combinations of $x_{\psi,1}, x_{\psi,2}, \ldots, x_{\psi,K-\eta}$ as received — up to noise level — at each user $k' \in [K]\backslash\psi$) which will help each user $k \in \psi$ resolve the already sent $x_{\psi,1}, x_{\psi,2}, \ldots, x_{\psi,K-\eta}$. This will be done in the next phase $j = \eta+2$.

[8]Whenever possible, we will henceforth avoid going into the details of the Q-MAT scheme. Some aspects of this scheme are similar to MAT, and a main new element is that Q-MAT applies digital transmission of interference, and a double-quantization method that collects and distributes residual interference across different rounds (this is here carried by $a_{k,t}^*$), in a manner that allows for ZF and MAT to coexist at maximal rates. Some of the details of this scheme are 'hidden' behind the choice of $\mathbf{G}_{c,t}$ and behind the loading of the MAT-type symbols $\boldsymbol{x}_{c,t}$ and additional auxiliary symbols $a_{k,t}^*$. The important element for the decoding part later on, will be how to load the symbols, the rate of each symbol, and the corresponding allocated power. An additional element that is hidden from the presentation here is that, while the Q-MAT scheme has many rounds, and while decoding spans more than one round, we will — in a slight abuse of notation — focus on describing just one round, which we believe is sufficient for the purposes of this paper here.

Phase $\eta+2$: The challenge now is for signals $\boldsymbol{x}_{c,t},\ t \in (T_{\eta+1}, T_{\eta+1}+T_{\eta+2}]$ to convey all the messages of the form

$$L_{\psi,k'},\ \forall k' \in [K]\backslash\psi,\ \forall \psi \in \Psi_{\eta+1}$$

to each receiver $k \in \psi$. Note that each of the above linear combinations, is now — during this phase — available (up to noise level) at the transmitter. Let

$$\Psi_{\eta+2} = \{\psi \in [K]\ :\ |\psi| = \eta+2\} \tag{44}$$

and consider for each $\psi \in \Psi_{\eta+2}$, a transmitted vector

$$\boldsymbol{x}_\psi = [x_{\psi,1}, \ldots, x_{\psi,K-\eta-1}, 0, \ldots, 0]^T$$

which carries the contents of $\eta+1$ different linear combinations $f_i(\{L_{\psi\backslash\{k\},k}\}_{k\in\psi}), i = 1,\ldots,\eta+1$ of the $\eta+2$ elements $\{L_{\psi\backslash\{k\},k}\}_{\forall k\in\psi}$ created by the transmitter. The linear combination coefficients defining each linear-combination function f_i, are predetermined and known at each receiver. The transmission of $\{\boldsymbol{x}_\psi\}_{\forall\psi\in\Psi_{\eta+2}}$ is sequential.

It is easy to see that there is a total of $(\eta+1)\binom{K}{\eta+2}$ symbols of the form $f_i(\{L_{\psi\backslash\{k\},k}\}_{k\in\psi}), i = 1,\ldots,\eta+1,\ \psi \in \Psi_{\eta+2}$, each of which can be considered as an order-$(\eta+2)$ signal intended for $\eta+2$ receivers in ψ. Using this, and following the same steps used in phase $\eta+1$, we calculate that

$$T_{\eta+2} = \binom{K}{\eta+2}\text{dur}(\boldsymbol{x}_\psi) = T_{\eta+1}\frac{\eta+1}{\eta+2}. \tag{45}$$

We now see that for each ψ, each receiver $k \in \psi$ recalls their own observation $L_{\psi\backslash\{k\},k}$ from the previous phase, and removes it from all the linear combinations $\{f_i(\{L_{\psi\backslash\{k\},k}\}_{\forall k\in\psi})\}_{i=1,\ldots,\eta+1}$, thus now being able to acquire the $\eta+1$ independent linear combinations $\{L_{\psi\backslash\{k'\},k'}\}_{\forall k'\in\psi\backslash\{k\}}$. The same holds for each other user $k' \in \psi$.

After this phase, we use $L_{\psi,k}, \psi \in \Psi_{\eta+2}$ to denote the received signal at receiver k. Like before, each receiver $k, k \in \psi$ needs $K-\eta-2$ extra observations of $x_{\psi,1},\ldots,x_{\psi,K-\eta-1}$ which will be seen from $L_{\psi,k'}, \forall k' \notin \psi$, which will come from order-$(\eta+3)$ messages that are created by the transmitter and which will be sent in the next phase.

Phase j $(\eta+3 \le j \le K)$: Generalizing the described approach to any phase $j \in [\eta+3,\ldots,K]$, we will use $\boldsymbol{x}_{c,t},\ t \in [\sum_{i=\eta+1}^{j-1} T_i, \sum_{i=\eta+1}^{j} T_i]$ to convey all the messages of the form

$$L_{\psi,k'},\ \forall k' \in [K]\backslash\psi,\ \forall \psi \in \Psi_{j-1}$$

to each receiver $k \in \psi$. For each

$$\psi \in \Psi_j := \{\psi \in [K]\ :\ |\psi| = j\} \tag{46}$$

each transmitted vector

$$\boldsymbol{x}_\psi = [x_{\psi,1}, \ldots, x_{\psi,K-j-1}, 0, \ldots, 0]^T$$

will carry the contents of $j-1$ different linear combinations $f_i(\{L_{\psi\backslash\{k\},k}\}_{k\in\psi}), i = 1,\ldots,j-1$ of the j elements $\{L_{\psi\backslash\{k\},k}\}_{\forall k\in\psi}$ created by the transmitter. After the sequential transmission of $\{\boldsymbol{x}_\psi\}_{\forall\psi\in\Psi_j}$, each receiver k can obtain the $j-1$ independent linear combinations $\{L_{\psi\backslash\{k'\},k'}\}_{\forall k'\in\psi\backslash\{k\}}$. The same holds for each other user $k' \in \psi$. As with the previous phases, we can see that

$$T_j = T_{\eta+1}\frac{\eta+1}{j},\ j = \eta+3,\ldots,K. \tag{47}$$

This process terminates with phase $j = K$, during which each

$$\boldsymbol{x}_\psi = [x_{\psi,1}, 0, 0, \ldots, 0]^T$$

carries a single scalar that is decoded easily by all. Based on this, backwards decoding will allow for users to retrieve $\{X_\psi\}_{\psi\in\Psi_{\eta+1}}$. This is described immediately afterwards. In treating the decoding part, we briefly recall that each $a_{k,t},\ k = 1,\ldots,K$ carries (during period $t \in [0,T]$), all of the uncached $W^{\mathrm{c}}_{R_k}$ and all of the unfolded $\{W^{c,\bar{f}}_{R_k,\tau}\}_{\tau\in\Psi_\eta\backslash\Psi^{(k)}_\eta}$.

D. Decoding

The whole transmission lasts $K-\eta$ phases. For each phase $j, j = \eta+1,\cdots,K$ and the corresponding ψ, the received signal $y_{k,t},\ t \in [\sum_{i=\eta+1}^{j-1} T_i, \sum_{i=\eta+1}^{j} T_i]$ of desired user k $(k \in \psi)$ takes the form

$$y_{k,t} = \underbrace{\overbrace{\boldsymbol{h}_{k,t}^T\mathbf{G}_{c,t}\boldsymbol{x}_{c,t}}^{\text{rate } 1-\alpha}}_{L_{\psi,k},\ \text{power } \doteq P} + \underbrace{\overbrace{\boldsymbol{h}_{k,t}^T\sum_{k\in\bar{\psi}}^{K}\boldsymbol{g}_{k,t}a^*_{k,t}}^{\text{rate } 1-\alpha}}_{\doteq P^{1-\alpha}} + \underbrace{\overbrace{\boldsymbol{h}_{k,t}^T\boldsymbol{g}_{k,t}a_{k,t}}^{\text{rate } \alpha}}_{P^\alpha} \tag{48}$$

The received signal $y_{k',t}$ of undesired user k' $(k' \in [K]\backslash\psi)$ takes the form

$$y_{k',t} = \underbrace{\overbrace{\boldsymbol{h}_{k',t}^T\sum_{k'\in\bar{\psi}}^{K}\boldsymbol{g}_{k',t}a^*_{k',t}}^{\text{rate } 1-\alpha}}_{\text{power } \doteq P} + \underbrace{\overbrace{\boldsymbol{h}_{k',t}^T\mathbf{G}_{c,t}\boldsymbol{x}_{c,t}}^{\text{rate } 1-\alpha}}_{L_{\psi,k'},\ \doteq P^{1-\alpha}} + \underbrace{\overbrace{\boldsymbol{h}_{k',t}^T\boldsymbol{g}_{k',t}a_{k',t}}^{\text{rate } \alpha}}_{P^\alpha} \tag{49}$$

where in both cases, we ignore the Gaussian noise and the ZF-related noise up to P^0. In addition to somehow send $L_{\psi,k'}$ to the next MAT phase, as we see in [48], after each phase, $L_{\psi,k'}$ is first quantized with $(1-2\alpha)\log P$ bits, which results in a residual quantizaton noise $n_{\psi,k'}$ with power scaling as P^α. Then, the transmitter quantizes the quantization noise $n_{\psi,k'}$ with an additional $\alpha\log P$ bits, which will be carried by the auxiliary data symbols $a^*_{k',t}$ in the corresponding phase in the next round (note: additional requests from the users). In this way, we can see that the 'common' signal $\boldsymbol{x}_{c,t}$ can be decoded with the assistance of an auxiliary data symbol from the next round. After this, each user k will remove $\boldsymbol{h}_{k,t}^T\mathbf{G}_{c,t}\boldsymbol{x}_{c,t}$ from their received signals, and readily decode their private symbols $a_{k,t},\ t \in [0,T]$, thus allowing for retrieval of their own unfolded $\{W^{c,\bar{f}}_{R_k,\psi\backslash\{k\}}\}_{\psi\in\Psi_{\eta+1}}$ and uncached $W^{\mathrm{c}}_{R_k}$. In terms of decoding the common information, as discussed above, each receiver k will perform a backwards reconstruction of the sets

of overheard equations

$$\begin{array}{c}\{L_{\psi,k'},\ \forall k' \in [K]\backslash\psi\}_{\forall\psi\in\Psi_K}\\ \downarrow\\ \{L_{\psi,k'},\ \forall k' \in [K]\backslash\psi\}_{\forall\psi\in\Psi_{K-1}}\\ \vdots\\ \downarrow\\ \{L_{\psi,k'},\ \forall k' \in [K]\backslash\psi\}_{\forall\psi\in\Psi_{\eta+2}}\end{array}$$

until phase $\eta+2$. At this point, each user k has enough observations to recover the original $K-\eta$ symbols $x_{\psi,1}, x_{\psi,2}, \ldots, x_{\psi,K-\eta}$ that fully convey X_ψ, hence each user k can reconstruct their own set $\{W^{c,f}_{R_k,\psi\backslash\{k\}}\}_{\psi\in\Psi_{\eta+1}}$ which, combined with the information from the $a_{k,t},\ t=[0,T]$ allow for each user k to reconstruct $\{W^{c}_{R_k,\psi\backslash\{k\}}\}_{\psi\in\Psi_{\eta+1}}$ which is then combined with Z_k to allow for reconstruction of the requested file W_{R_k}.

E. Calculation of T

To calculate T, we recall from (47) that

$$\begin{aligned}T &= \sum_{j=\eta+1}^{K} T_j = T_{\eta+1}\sum_{j=\eta+1}^{K}\frac{\eta+1}{j}\\ &= (\eta+1)(H_K-H_\eta)T_{\eta+1}\end{aligned} \tag{50}$$

which combines with (41) and (43) to give

$$T = \frac{(K-\Gamma)(H_K-H_\eta)}{(K-\eta)+\alpha(\eta+K(H_K-H_\eta-1))} \tag{51}$$

as stated in Theorem 1. The bound by $T=1-\gamma$ seen in the theorem, corresponds to the fact that the above expression (51) applies, as is, only when $\alpha \le \alpha_{b,K-1} = \frac{K(1-\gamma)-1}{(K-1)(1-\gamma)}$ which corresponds to $\eta = K-1$ (where X_ψ are fully common messages, directly desired by all), for which we already get the best possible $T=1-\gamma$, and hence there is no need to go beyond $\alpha = \frac{K(1-\gamma)-1}{(K-1)(1-\gamma)}$.

V. Conclusions

This work studied the previously unexplored interplay between coded-caching and CSIT feedback quality and timeliness. This is motivated by the fact that CSIT and coded caching are two powerful ingredients that are hard to obtain, and by the fact that these ingredients are intertwined in a synergistic and competing manner. In addition to the substantial cache-aided DoF gains revealed here, the results suggest the interesting practical ramification that distributing predicted content 'during the night', can offer continuous amelioration of the load of predicting and disseminating CSIT during the day.

VI. Appendix

A. Lower bound on T^* (proof of Lemma 1)

This part draws from the bound in [1], and it is similar to that in [49] which deals with the case of $\alpha=0$. To lower bound T, we consider the easier problem where we want to serve $s \le K$ different files to s users, each with access to all caches. We also consider that we repeat this (easier) last experiment $\lfloor \frac{N}{s} \rfloor$ times, thus spanning a total duration of $T\lfloor \frac{N}{s} \rfloor$ (and up to $\lfloor \frac{N}{s} \rfloor s$ files delivered). At this point, we transfer to the equivalent setting of the s-user MISO BC with delayed CSIT and imperfect current CSIT, and a side-information multicasting link to the receivers, of capacity d_m (files per time slot). Under the assumption that in this latter setting, decoding happens at the end of communication, and once we set

$$d_m T\lfloor \frac{N}{s} \rfloor = sM \tag{52}$$

(which guarantees that the side information from the side link, throughout the communication process, matches the maximum amount of information in the caches), we have that

$$T\lfloor \frac{N}{s} \rfloor d'_\Sigma(d_m) \ge \lfloor \frac{N}{s} \rfloor s \tag{53}$$

where $d'_\Sigma(d_m)$ is any sum-DoF upper bound on the above s-user MISO BC channel with delayed CSIT and the aforementioned side link. Using the bound

$$d'_\Sigma(d_m) = s\alpha + \frac{s}{H_s}(1-\alpha+d_m)$$

from Lemma 2, and applying (52), we get

$$T\lfloor \frac{N}{s} \rfloor \big(s\alpha + \frac{s}{H_s}(1-\alpha+\frac{sM}{T\lfloor \frac{N}{s} \rfloor})\big) \ge \lfloor \frac{N}{s} \rfloor s \tag{54}$$

and thus that

$$T \ge \frac{1}{(H_s\alpha+1-\alpha)}(H_s - \frac{Ms}{\lfloor \frac{N}{s} \rfloor}) \tag{55}$$

which implies a lower bound on the original s-user problem. Maximization over all s, gives the desired bound on the optimal T^*

$$T^* \ge \max_{s\in\{1,\ldots,\lfloor \frac{N}{M} \rfloor\}} \frac{1}{(H_s\alpha+1-\alpha)}(H_s - \frac{Ms}{\lfloor \frac{N}{s} \rfloor}) \tag{56}$$

required for the original K-user problem. This concludes the proof.

B. Bounding the sum-DoF of the s-user MISO BC, with delayed CSIT, α-quality current CSIT, and additional side information

We begin with the statement of the lemma, which we prove immediately below.

Lemma 2: For the s-user MISO BC, with delayed CSIT, α-quality current CSIT, and an additional parallel side-link of capacity that scales as $d_m \log P$, the sum-DoF is upper bounded as

$$d_\Sigma(d_m) \le s\alpha + \frac{s}{H_s}(1-\alpha+d_m). \tag{57}$$

Proof: Our proof traces the proof of [51], adapting for the additional α-quality current CSIT.

Consider a permutation π of the set $\mathcal{E}=\{1,2,\cdots,s\}$. For any user $k, k\in\mathcal{E}$, we provide the received signals $y_k^{[n]}$ as well as the message W_k of user k to user $k+1, k+2, \cdots, s$. We define the following notation

$$\Omega^{[n]} := \{\boldsymbol{h}_k^{[n]}\}_{k=1}^s,\ \hat{\Omega}^{[n]} := \{\hat{\boldsymbol{h}}_k^{[n]}\}_{k=1}^s,\ \mathcal{U}^{[n]} := \{\Omega^{[n]}, \hat{\Omega}^{[n]}\},$$

$$\boldsymbol{h}_k^{[t]} := \{\boldsymbol{h}_k^{(i)}\}_{i=1}^t, \quad y_k^{[t]} := \{y_k^{(i)}\}_{i=1}^t, t = 1, 2, \cdots, n,$$

$$W_{[k]} := \{W_1, W_2, \cdots, W_k\}, \quad y_{[k]}^{[n]} := \{y_1^{[n]}, y_2^{[n]}, \cdots, y_k^{[n]}\}.$$

Then for $k = 1, 2, \cdots, s$, we have

$$\begin{aligned}
&n(R_k - \epsilon_n) \\
&\le I(W_k; y_{[k]}^{[n]}, y_0^{[n]}, W_{[k-1]}|\mathcal{U}^{[n]}) && (58)\\
&= I(W_k; y_{[k]}^{[n]}, y_0^{[n]}|W_{[k-1]}, \mathcal{U}^{[n]}) && (59)\\
&= I(W_k; y_{[k]}^{[n]}|W_{[k-1]}, \mathcal{U}^{[n]}) + I(W_k; y_0^{[n]}|y_{[k]}^{[n]}, W_{[k-1]}, \mathcal{U}^{[n]}) \\
&= h(y_{[k]}^{[n]}|W_{[k-1]}, \mathcal{U}^{[n]}) - h(y_{[k]}^{[n]}|W_{[k]}, \mathcal{U}^{[n]}) \\
&\quad + h(y_0^{[n]}|y_{[k]}^{[n]}, W_{[k-1]}, \mathcal{U}^{[n]}) - h(y_0^{[n]}|y_{[k]}^{[n]}, W_{[k]}, \mathcal{U}^{[n]}) && (60)
\end{aligned}$$

where (58) follows from Fano's inequality, where (59) holds due to the fact that the messages are independent, and where the last two steps use the basic chain rule. Note that $W_0 = 0$.

$$\begin{aligned}
&\sum_{k=1}^{s-1}\Big(\frac{h(y_{[k+1]}^{[n]}|W_{[k]}, \mathcal{U}^{[n]})}{k+1} - \frac{h(y_{[k]}^{[n]}|W_{[k]}, \mathcal{U}^{[n]})}{k}\Big) \\
&= \sum_{t=1}^{[n]}\sum_{k=1}^{s-1}\Big(\frac{h(y_1^{(t)}, \cdots, y_{k+1}^{(t)}|y_1^{[t-1]}, \cdots, y_{k+1}^{[t-1]}, W_{[k]}, \mathcal{U}^{[n]})}{k+1} \\
&\quad - \frac{h(y_1^{(t)}, \cdots, y_k^{(t)}|y_1^{[t-1]}, \cdots, y_k^{[t-1]}, W_{[k]}, \mathcal{U}^{[n]})}{k}\Big) && (61)\\
&= \sum_{t=1}^{[n]}\sum_{k=1}^{s-1}\Big(\frac{h(y_1^{(t)}, \cdots, y_{k+1}^{(t)}|y_1^{[t-1]}, \cdots, y_{k+1}^{[t-1]}, W_{[k]}, \mathcal{U}^{[t]})}{k+1} \\
&\quad - \frac{h(y_1^{(t)}, \cdots, y_k^{(t)}|y_1^{[t-1]}, \cdots, y_k^{[t-1]}, W_{[k]}, \mathcal{U}^{[t]})}{k}\Big) && (62)\\
&\le \sum_{t=1}^{[n]}\sum_{k=1}^{s-1}\Big(\frac{h(y_1^{(t)}, \cdots, y_{k+1}^{(t)}|y_1^{[t-1]}, \cdots, y_{k+1}^{[t-1]}, W_{[k]}, \mathcal{U}^{[t]})}{k+1} \\
&\quad - \frac{h(y_1^{(t)}, \cdots, y_k^{(t)}|y_1^{[t-1]}, \cdots, y_{k+1}^{[t-1]}, W_{[k]}, \mathcal{U}^{[t]})}{k}\Big) && (63)\\
&\le \sum_{t=1}^{[n]}\sum_{k=1}^{s-1}\frac{1}{k+1}\alpha\log P + n\cdot o(\log P) && (64)\\
&\le n(H_s - 1)\alpha\log P + n\cdot o(\log P) && (65)
\end{aligned}$$

where (61) follows from the linearity of the summation, where (62) holds since the received signal is independent of the future channel state information, where (63) uses the fact that conditioning reduces entropy, and where (64) is from the fact that Gaussian distribution maximizes differential entropy under the covariance constraint and from [9, Lemma 2]. From (60), we then have

$$\begin{aligned}
&\sum_{k=1}^{s}\frac{n(R_k - \epsilon_n)}{k} \\
&\le \sum_{k=1}^{s-1}\Big(\frac{h(y_{[k+1]}^{[n]}|W_{[k]}, \mathcal{U}^{[n]})}{k+1} - \frac{h(y_{[k]}^{[n]}|W_{[k]}, \mathcal{U}^{[n]})}{k}\Big) \\
&\quad + h(y_1^{[n]}|\mathcal{U}^{[n]}) - \frac{1}{s}h(y_{[s]}^{[n]}|W_{[s]}, \mathcal{U}^{[n]}) \\
&\quad + \sum_{k=1}^{s-1}\Big(\frac{H(y_0^{[n]}|y_{[k+1]}^{[n]}, W_{[k]}, \mathcal{U}^{[n]})}{k+1} - \frac{H(y_0^{[n]}|y_{[k]}^{[n]}, W_{[k]}, \mathcal{U}^{[n]})}{k}\Big) \\
&\quad + H(y_0^{[n]}|y_1^{[n]}, \mathcal{U}^{[n]}) - \frac{1}{s}H(y_0^{[n]}|y_{[s]}^{[n]}, W_{[s]}, \mathcal{U}^{[n]}) \\
&\le n(H_s - 1)\alpha\log P + \underbrace{h(y_1^{[n]}|\mathcal{U}^{[n]})}_{\le n\log P} + \underbrace{H(y_0^{[n]}|y_1^{[n]}, \mathcal{U}^{[n]})}_{\le n\cdot d_m\log P} \\
&\quad + \sum_{k=1}^{s-1}\Big((\frac{1}{k+1} - \frac{1}{k})H(y_0^{[n]}|y_{[k]}^{[n]}, W_{[k]}, \mathcal{U}^{[n]}\Big) + n\cdot o(\log P) \\
&\le n(H_s - 1)\alpha\log P + n\log P + n\cdot d_m\log P + n\cdot o(\log P). && (66)
\end{aligned}$$

Dividing by $n\log P$ and letting $P \to \infty$ gives

$$\sum_{k=1}^{s}\frac{d_k}{k} \le (H_s - 1)\alpha + 1 + d_m \tag{67}$$

which implies that

$$d_\Sigma(d_m) \le s\alpha + \frac{s}{H_s}(1 - \alpha + d_m) \tag{68}$$

which completes the proof of Lemma 2.

■

C. Bounding the gap between the achievable T and the optimal T^*

Our aim here is to show that

$$\frac{T(\gamma, \alpha > 0)}{T^*(\gamma, \alpha > 0)} < 4$$

and we will do so by showing that the above gap is smaller than the gap we calculated in [49] for $\alpha = 0$, which was again bounded above by 4. For this, we will use the expression[9]

$$T(\gamma, \alpha > 0) = \frac{(1-\gamma)(H_K - H_\Gamma)}{\alpha(H_K - H_\Gamma) + (1-\alpha)(1-\gamma)} \tag{69}$$

from Theorem 2, and the expression

$$T^*(\gamma, \alpha > 0) \ge \max_{s\in\{1,\ldots,\lfloor\frac{N}{M}\rfloor\}}\frac{1}{(H_s\alpha + 1 - \alpha)}(H_s - \frac{Ms}{\lfloor\frac{N}{s}\rfloor}) \tag{70}$$

from Lemma 1. Hence we have

$$\begin{aligned}
\frac{T}{T^*} &\le \frac{\frac{(1-\gamma)(H_K - H_{K\gamma})}{\alpha(H_K - H_{K\gamma}) + (1-\alpha)(1-\gamma)}}{\max_{s\in\{1,\ldots,\lfloor\frac{N}{M}\rfloor\}}\frac{1}{(H_s\alpha+1-\alpha)}(H_s - \frac{Ms}{\lfloor\frac{N}{s}\rfloor})} && (71)\\
&\le \underbrace{\frac{\frac{(1-\gamma)(H_K - H_{K\gamma})}{\alpha(H_K - H_{K\gamma}) + (1-\alpha)(1-\gamma)}}{\frac{1}{(H_{s_c}\alpha+1-\alpha)}(H_{s_c} - \frac{Ms_c}{\lfloor\frac{N}{s_c}\rfloor})}}_{g(s_c,\gamma)} && (72)
\end{aligned}$$

where $s = s_c \in \{1, \ldots, \lfloor\frac{N}{M}\rfloor\}$, but where this s_c will be chosen here to be exactly the same as in the case of $\alpha = 0$.

[9]We note that the here derived upper bound on the gap corresponding to the T in Theorem 2, automatically applies as an upper bound to the gap corresponding to the T from Theorem 1, because the latter T is smaller than the former.

This will be useful because, for that case of $\alpha = 0$, we have already proved that the same specific s_c guarantees that

$$\frac{H_K - H_{K\gamma}}{H_{s_c} - \frac{Ms_c}{\lfloor \frac{N}{s_c} \rfloor}} < 4 \tag{73}$$

for the appropriate ranges of γ. This will apply towards bounding (72).

The proof is broken in two cases, corresponding to $\gamma \in [\frac{1}{36}, \frac{K-1}{K}]$, and $\gamma \in [0, \frac{1}{36}]$.

1) Case 1 ($\alpha > 0, \gamma \in [\frac{1}{36}, \frac{K-1}{K}]$): As when $\alpha = 0$ (cf. [49]), we again set $s = 1$, which reduces (72) to

$$\frac{T(\alpha > 0, \gamma)}{T^*(\alpha > 0, \gamma)} \leq \frac{\frac{(1-\gamma)(H_K - H_{K\gamma})}{\alpha(H_K - H_{K\gamma}) + (1-\alpha)(1-\gamma)}}{1-\gamma}.$$

For this case — when α was zero, and when we chose the same $s = 1$ — we have already proved that

$$\frac{T(\alpha = 0, \gamma)}{1-\gamma} < 4.$$

As a result, since $T(\alpha > 0, \gamma) < T(\alpha = 0, \gamma)$, and since $1 - \gamma \leq T^*$, we conclude that

$$\frac{T(\alpha > 0, \gamma)}{T^*} < 4, \ \gamma \in [\frac{1}{36}, \frac{K-1}{K}]$$

which completes this part of the proof.

2) Case 2 ($\alpha > 0, \gamma \in [0, \frac{1}{36}]$): Going back to (72), we now aim to bound

$$g(s_c, \gamma) := \frac{\frac{(1-\gamma)(H_K - H_{K\gamma})}{\alpha(H_K - H_{K\gamma}) + (1-\alpha)(1-\gamma)}}{\frac{1}{(H_{s_c}\alpha + 1 - \alpha)}(H_{s_c} - \frac{Ms_c}{\lfloor \frac{N}{s_c} \rfloor})} < 4. \tag{74}$$

We already know from the case of $\alpha = 0$ (cf. (73)) that

$$\frac{H_K - H_{K\gamma}}{H_{s_c} - \frac{Ms_c}{\lfloor \frac{N}{s_c} \rfloor}} < 4 \tag{75}$$

holds. Hence we will prove that

$$g(s_c, \gamma) \leq \frac{H_K - H_{K\gamma}}{H_{s_c} - \frac{Ms_c}{\lfloor \frac{N}{s_c} \rfloor}} \tag{76}$$

to guarantee the bound. We note that (76) is implied by

$$H_{s_c} \leq \frac{H_K - H_{K\gamma}}{1-\gamma} \tag{77}$$

which is implied by

$$\log(s_c) \leq \frac{\log(1/\gamma)}{1-\gamma} - \epsilon_6, \ \epsilon_6 = H_6 - \log(6) \tag{78}$$

because $H_{s_c} \leq \log(s_c) + \epsilon_6, \forall s_c \geq 6, \forall \gamma \in [0, \frac{1}{36}], \forall K$. Furthermore (78) is implied by

$$\frac{1}{2}\log(\frac{1}{\gamma}) \leq \frac{\log(1/\gamma)}{1-\gamma} - \epsilon_6 \tag{79}$$

because $\gamma \in [\frac{1}{(s_c+1)^2}, \frac{1}{s_c^2}]$ means that $s_c \leq \sqrt{\frac{1}{\gamma}}$. Since $\frac{1}{1-\gamma} \geq 1$, then (79) is implied by

$$\frac{1}{2}\log(\frac{1}{\gamma}) \leq \log(\frac{1}{\gamma}) - \epsilon_6. \tag{80}$$

It is obvious that (80) holds since $\gamma \leq \frac{1}{36}$. Towards this, by proving (80), we guarantee (74) and the desired bound. This completes the proof.

REFERENCES

[1] M. Maddah-Ali and U. Niesen, "Fundamental limits of caching," *Information Theory, IEEE Transactions on*, vol. 60, no. 5, pp. 2856–2867, May 2014.

[2] S. Yang, M. Kobayashi, D. Gesbert, and X. Yi, "Degrees of freedom of time correlated MISO broadcast channel with delayed CSIT," *IEEE Trans. Inf. Theory*, vol. 59, no. 1, pp. 315–328, Jan. 2013.

[3] J. Chen and P. Elia, "Toward the performance versus feedback tradeoff for the two-user miso broadcast channel," *IEEE Trans. Inf. Theory*, vol. 59, no. 12, pp. 8336–8356, Dec. 2013.

[4] A. G. Davoodi and S. A. Jafar, "Aligned image sets under channel uncertainty: Settling a conjecture by Lapidoth, Shamai and Wigger on the collapse of degrees of freedom under finite precision CSIT," *CoRR*, vol. abs/1403.1541, 2014. [Online]. Available: http://arxiv.org/abs/1403.1541

[5] G. Caire, N. Jindal, M. Kobayashi, and N. Ravindran, "Multiuser MIMO achievable rates with downlink training and channel state feedback," *IEEE Trans. Inf. Theory*, vol. 56, no. 6, pp. 2845 – 2866, Jun. 2010.

[6] M. A. Maddah-Ali and D. N. C. Tse, "Completely stale transmitter channel state information is still very useful," *IEEE Trans. Inf. Theory*, vol. 58, no. 7, pp. 4418 – 4431, Jul. 2012.

[7] T. Gou and S. Jafar, "Optimal use of current and outdated channel state information: Degrees of freedom of the MISO BC with mixed CSIT," *IEEE Communications Letters*, vol. 16, no. 7, pp. 1084 – 1087, Jul. 2012.

[8] J. Chen and P. Elia, "Degrees-of-freedom region of the MISO broadcast channel with general mixed-CSIT," in *Proc. Information Theory and Applications Workshop (ITA)*, Feb. 2013.

[9] P. de Kerret, X. Yi, and D. Gesbert, "On the degrees of freedom of the K-user time correlated broadcast channel with delayed CSIT," Jan. 2013, available on arXiv:1301.2138.

[10] J. Chen, S. Yang, and P. Elia, "On the fundamental feedback-vs-performance tradeoff over the MISO-BC with imperfect and delayed CSIT," Jul. 2013, in *ISIT13*, available on arXiv:1302.0806.

[11] C. Vaze and M. Varanasi, "The degree-of-freedom regions of MIMO broadcast, interference, and cognitive radio channels with no CSIT," *IEEE Trans. Inf. Theory*, vol. 58, no. 8, pp. 5254 – 5374, Aug. 2012.

[12] R. Tandon, S. A. Jafar, S. Shamai, and H. V. Poor, "On the synergistic benefits of alternating CSIT for the MISO BC," *IEEE Trans. Inf. Theory*, vol. 59, no. 7, pp. 4106 – 4128, Jul. 2013.

[13] N. Lee and R. W. Heath Jr., "Not too delayed CSIT achieves the optimal degrees of freedom," in *Proc. Allerton Conf. Communication, Control and Computing*, Oct. 2012.

[14] C. Hao and B. Clerckx, "Imperfect and unmatched CSIT is still useful for the frequency correlated MISO broadcast channel," in *Proc. IEEE Int. Conf. Communications (ICC)*, Budapest, Hungary, Jun. 2013.

[15] C. S. Vaze and M. K. Varanasi, "The degrees of freedom region of two-user and certain three-user MIMO broadcast channel with delayed CSI," Dec. 2011, submitted to *IEEE Trans. Inf. Theory*, available on arXiv:1101.0306.

[16] M. J. Abdoli, A. Ghasemi, and A. K. Khandani, "On the degrees of freedom of three-user MIMO broadcast channel with delayed CSIT," in *Proc. IEEE Int. Symp. Information Theory (ISIT)*, Jul. 2011.

[17] N. Lee and R. W. Heath, "Not too delayed CSIT achieves the optimal degrees of freedom," in *Proc. Allerton Conference on Communication, Control, and Computing (Allerton)*, 2012.

[18] R. Tandon, S. A. Jafar, S. Shamai (Shitz), and H. V. Poor, "On the synergistic benefits of alternating CSIT for the MISO BC," *IEEE Trans. Inf. Theory*, vol. 59, no. 7, pp. 4106–4128, 2013.

[19] M. Torrellas, A. Agustin, and J. Vidal, "Retrospective interference alignment for the MIMO interference broadcast channel," in *Proc. IEEE International Symposium on Information Theory (ISIT)*, June 2015.

[20] A. Bracher and M. A. Wigger, "Feedback and partial message side-information on the semideterministic broadcast channel," in *Proc. IEEE International Symposium on Information Theory (ISIT)*, June 2015.

[21] S. Lashgari, R. Tandon, and S. Avestimehr, "Three-user MISO broadcast channel: How much can CSIT heterogeneity help?" in *Proc. IEEE International Conference on Communications (ICC)*, June 2015, pp. 4187–4192.

[22] M. R. Korupolu, C. G. Plaxton, and R. Rajaraman, "Placement algorithms for hierarchical cooperative caching," in *Proc. ACM-SIAM SODA*, Jan. 1999, pp. 586–595.

[23] B.-J. Ko and D. Rubenstein, "Distributed self-stabilizing placement of replicated resources in emerging networks," *IEEE/ACM Trans. Netw.*, vol. 13, no. 3, pp. 476–487, Jun. 2005.

[24] Y. Birk and T. Kol, "Coding on demand by an informed source (ISCOD) for efficient broadcast of different supplemental data to caching clients," *IEEE Trans. Inf. Theory*, vol. 52, no. 6, pp. 2825–2830, Jun. 2006.

[25] S. Borst, V. Gupta, and A. Walid, "Distributed caching algorithms for content distribution networks," in *INFOCOM, 2010 Proceedings IEEE*, Mar. 2010, pp. 1–9.

[26] S. Wang, W. Li, X. Tian, and H. Liu, "Fundamental limits of heterogenous cache," *CoRR*, vol. abs/1504.01123, 2015. [Online]. Available: http://arxiv.org/abs/1504.01123

[27] M. A. Maddah-Ali and U. Niesen, "Decentralized caching attains order-optimal memory-rate tradeoff," *CoRR*, vol. abs/1301.5848, 2013. [Online]. Available: http://arxiv.org/abs/1301.5848

[28] M. Ji, A. M. Tulino, J. Llorca, and G. Caire, "Order optimal coded delivery and caching: Multiple groupcast index coding," *CoRR*, vol. abs/1402.4572, 2014. [Online]. Available: http://arxiv.org/abs/1402.4572

[29] H. Ghasemi and A. Ramamoorthy, "Improved lower bounds for coded caching," *CoRR*, vol. abs/1501.06003, 2015. [Online]. Available: http://arxiv.org/abs/1501.06003

[30] C. Wang, S. H. Lim, and M. Gastpar, "Information-theoretic caching: Sequential coding for computing," *CoRR*, vol. abs/1504.00553, 2015. [Online]. Available: http://arxiv.org/abs/1504.00553

[31] A. N., N. S. Prem, V. M. Prabhakaran, and R. Vaze, "Critical database size for effective caching," *CoRR*, vol. abs/1501.02549, 2015. [Online]. Available: http://arxiv.org/abs/1501.02549

[32] W. Huang, S. Wang, L. Ding, F. Yang, and W. Zhang, "The performance analysis of coded cache in wireless fading channel," *CoRR*, vol. abs/1504.01452, 2015. [Online]. Available: http://arxiv.org/abs/1504.01452

[33] R. Timo and M. A. Wigger, "Joint cache-channel coding over erasure broadcast channels," *CoRR*, vol. abs/1505.01016, 2015. [Online]. Available: http://arxiv.org/abs/1505.01016

[34] M. A. Maddah-Ali and U. Niesen, "Cache-aided interference channels," in *Proceedings of the IEEE International Symposium on Information Theory (ISIT'2015)*, Hong-Kong, China, 2015.

[35] N. Golrezaei, K. Shanmugam, A. Dimakis, A. Molisch, and G. Caire, "Femtocaching: Wireless video content delivery through distributed caching helpers," in *INFOCOM, 2012 Proceedings IEEE*, March 2012, pp. 1107–1115.

[36] B. Perabathini, E. Bastug, M. Kountouris, M. Debbah, and A. Conte, "Caching at the edge: a green perspective for 5G networks," *CoRR*, vol. abs/1503.05365, 2015. [Online]. Available: http://arxiv.org/abs/1503.05365

[37] M. Ji, M. F. Wong, A. Tulino, J. Llorca, G. Caire, M. Effros, and M. Langberg, "On the fundamental limits of caching in combination networks," in *Signal Processing Advances in Wireless Communications (SPAWC), 2015 IEEE 16th International Workshop on*, June 2015, pp. 695–699.

[38] U. Niesen, D. Shah, and G. W. Wornell, "Caching in wireless networks," *IEEE Trans. Inf. Theory*, vol. 58, no. 10, pp. 6524–6540, Oct 2012.

[39] E. Bastug, M. Bennis, and M. Debbah, "A transfer learning approach for cache-enabled wireless networks," *CoRR*, vol. abs/1503.05448, 2015. [Online]. Available: http://arxiv.org/abs/1503.05448

[40] A. F. Molisch, G. Caire, D. Ott, J. R. Foerster, D. Bethanabhotla, and M. Ji, "Caching eliminates the wireless bottleneck in video-aware wireless networks," *CoRR*, vol. abs/1405.5864, 2014. [Online]. Available: http://arxiv.org/abs/1405.5864

[41] J. Hachem, N. Karamchandani, and S. N. Diggavi, "Coded caching for heterogeneous wireless networks with multi-level access," *CoRR*, vol. abs/1404.6560, 2014. [Online]. Available: http://arxiv.org/abs/1404.6560

[42] J. Hachem, N. Karamchandani, and S. Diggavi, "Effect of number of users in multi-level coded caching," in *Proc. IEEE Int. Symp. Information Theory (ISIT)*, Hong-Kong, China, 2015.

[43] K. Shanmugam, M. Ji, A. Tulino, J. Llorca, and A. Dimakis, "Finite length analysis of caching-aided coded multicasting," 2015, submitted to *IEEE Trans. Inform. Theory - July 2015*.

[44] M. Deghel, E. Bastug, M. Assaad, and M. Debbah, "On the benefits of edge caching for MIMO interference alignment," in *Signal Processing Advances in Wireless Communications (SPAWC), 2015 IEEE 16th International Workshop on*, June 2015, pp. 655–659.

[45] A. Ghorbel, M. Kobayashi, and S. Yang, "Cache-enabled broadcast packet erasure channels with state feedback," *CoRR*, vol. abs/1509.02074, 2015. [Online]. Available: http://arxiv.org/abs/1509.02074

[46] M. Gatzianas, L. Georgiadis, and L. Tassiulas, "Multiuser broadcast erasure channel with feedback — capacity and algorithms," *IEEE Trans. Inf. Theory*, vol. 59, no. 9, pp. 5779–5804, Sept 2013.

[47] J. Zhang, F. Engelmann, and P. Elia, "Coded caching for reducing CSIT-feedback in wireless communications," in *Proc. Allerton Conf. Communication, Control and Computing*, Monticello, Illinois, USA, Sep. 2015.

[48] P. de Kerret, D. Gesbert, J. Zhang, and P. Elia, "Optimal sum-DoF of the K-user MISO BC with current and delayed feedback," 2016. [Online]. Available: https://arxiv.org/abs/1604.01653

[49] J. Zhang and P. Elia, "The synergistic gains of coded caching and delayed feedback," *CoRR*, vol. abs/1604.06531, Apr. 2016. [Online]. Available: http://arxiv.org/abs/1604.06531

[50] M. Kobayashi and G. Caire, "On the net DoF comparison between ZF and MAT over time-varying MISO broadcast channels," in *Proc. IEEE Int. Symp. Information Theory (ISIT)*, Jul. 2012.

[51] J. Chen, S. Yang, A. Özgür, and A. Goldsmith, "Achieving full dof in heterogeneous parallel broadcast channels with outdated CSIT," *CoRR*, vol. abs/1409.6808, 2014. [Online]. Available: http://arxiv.org/abs/1409.6808

[52] J. Zhang and P. Elia, "Fundamental limits of cache-aided wireless BC: Interplay of coded-caching and CSIT feedback," August 25 2015, *EURECOM report No. RR-15-307*, available on: http://www.eurecom.fr/publication/4723.